Contribution To: Open Problems Presented at the 2012 Canadian Conference on Computational Geometry

Joe Mitchell (Session Chair/Scribe)

August 8, 2012

On Wednesday afternoon, August 8, 2012, we held an open problem discussion at CCCG, Charlottetown, PEI. The problems span a range of topics, including fundamental algorithms, discrete geometry, algebra, combinatorics, and optimization.

I (Scott A. Mitchell) described this problem and submitted the following writeup to the open problem forum. I'll probably remove this local version when all the open problems are online.

Problem 1 (Scott A. Mitchell)**.** *Characterize the output of the Poisson-Disk process:* (This might be considered a problem in spatial statistics, but there are ties to Delaunay refinement and sphere packings.) Maximal Poisson-disk sampling (MPS) is a particular statistical process for generating a point cloud. The location for the next point is chosen uniformly by area at random. A point has an empty disk of radius r around it; if a new point falls into a prior point's disk, it is rejected and not added to the sample. The process continues until the sampling is maximal: the entire domain is covered by samples' disks and there is no room for another sample. Let the domain be a two-dimensional square with periodic (toroidal) boundary conditions, so there are no domain boundary issues to consider. A math definition appears in [Efficient-MPS]:

`http://www.cs.sandia.gov/~samitch/bibliography_2007.html#efficient-mps`

I am aware of no analytic description of what the correct output of MPS is supposed to be. I haven't even seen an experimental characterization! As such, currently for an algorithm to be correct, it must be step-by-step equivalent to the statistical process. For an example algorithm like this, see again [Efficient-MPS]. A characterization of the output is important because it would enable the design of more efficient algorithms. A metaphor is that bubble-sort is a process, but the characterization of its output as "sorted order" allows the discovery of e.g. quicksort to generate "sorted order" more efficiently.

The computer graphics community typically measures the output of MPS by generating Fourier transform pictures of the output. See "Point Set Analysis" [PSA], for software and paper references for a standard way of generating these pictures:

`http://code.google.com/p/psa/`

My understanding of PSA follows. The vectors of distances between all pairs of points are calculated. The Fourier transform of the distance vectors are taken and displayed, and a picture with oscillating dark and light rings is expected. Integrating this transform over concentric circles produces a one-dimensional graph by increasing radius. (A nuance is how to bin distances to generate smooth pictures.) Figure 1 top shows the kinds of pictures the Graphics community expects to see for MPS.

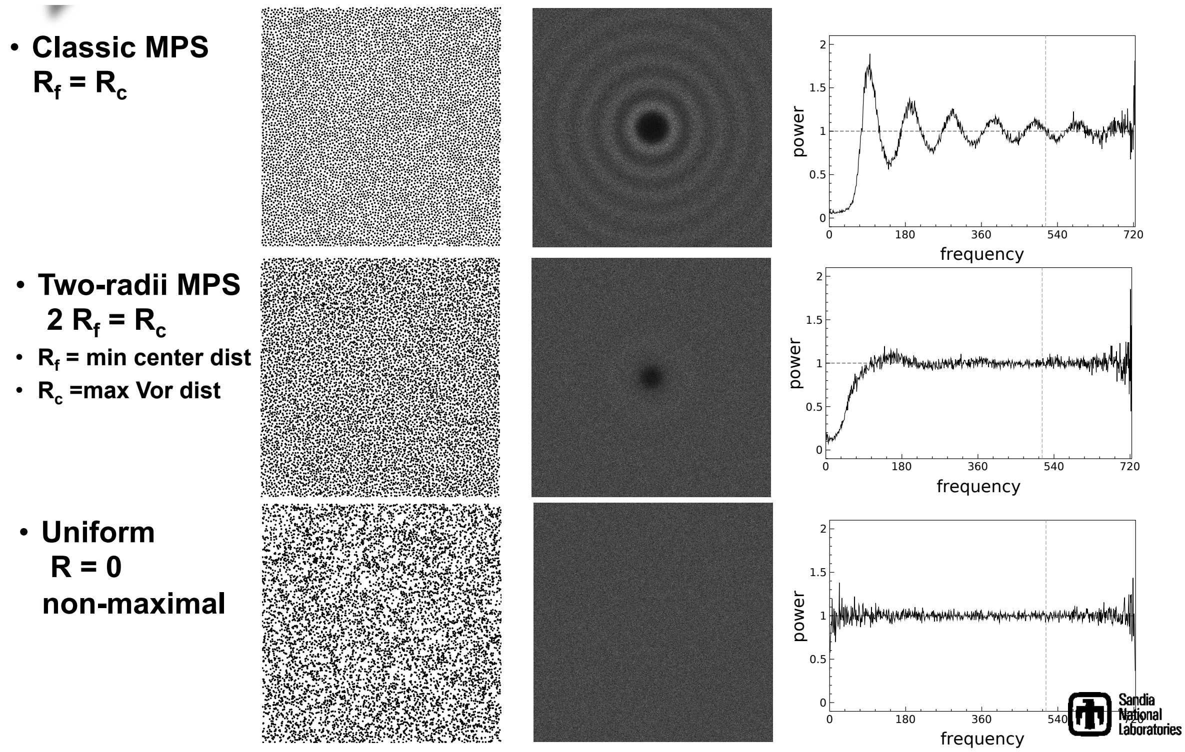

Figure 1: Point clouds visualized using PSA. Top is standard MPS, and middle two-radii MPS from CCCG 2012. The bottom is a uniform random point cloud without inhibition disks, using about the same number of points.

Subproblem A: Can you characterize the PSA pictures for MPS, especially Figure 1 top right? What is the mean location and height of the peaks? What is their standard deviation? Is the distribution around the mean normal? (Recall MPS is a random process.) Perhaps an experimental characterization is an easier place to start than an analytic characterization.

Subproblem B: Is some variant of MPS better than standard MPS for texture synthesis graphics applications? At CCCG 2012 I presented a paper "Variable Radii MPS." The two-radii MPS variant generates a spectrum with less oscillations; see Figure 1 middle. We suspect, but don't know for sure, if this is better for applications.

MPS produces a sphere packing, halve the disk radii r then the disks do not overlap. This is a well-spaced point set. Delaunay refinement also produces a well-spaced point set. Sometimes the PSA pictures of the output of Delaunay refinement look similar to MPS, sometimes not, depending on the target edge length, angle threshold, the use of off-centers, etc.; see Figure 2.

Subproblem C: characterize the PSA pictures (Fourier spectrum) of the output of Delaunay refinement and its variants.

In computational geometry we often measure point sets by the angles and edge length histograms in a Delaunay triangulation of the points. These histograms are different for MPS point clouds than for Delaunay refinement output; see Figure 3.

Subproblem D: Characterize MPS output using computational geometry measures of Delaunay triangulation edge lengths and angle distributions.

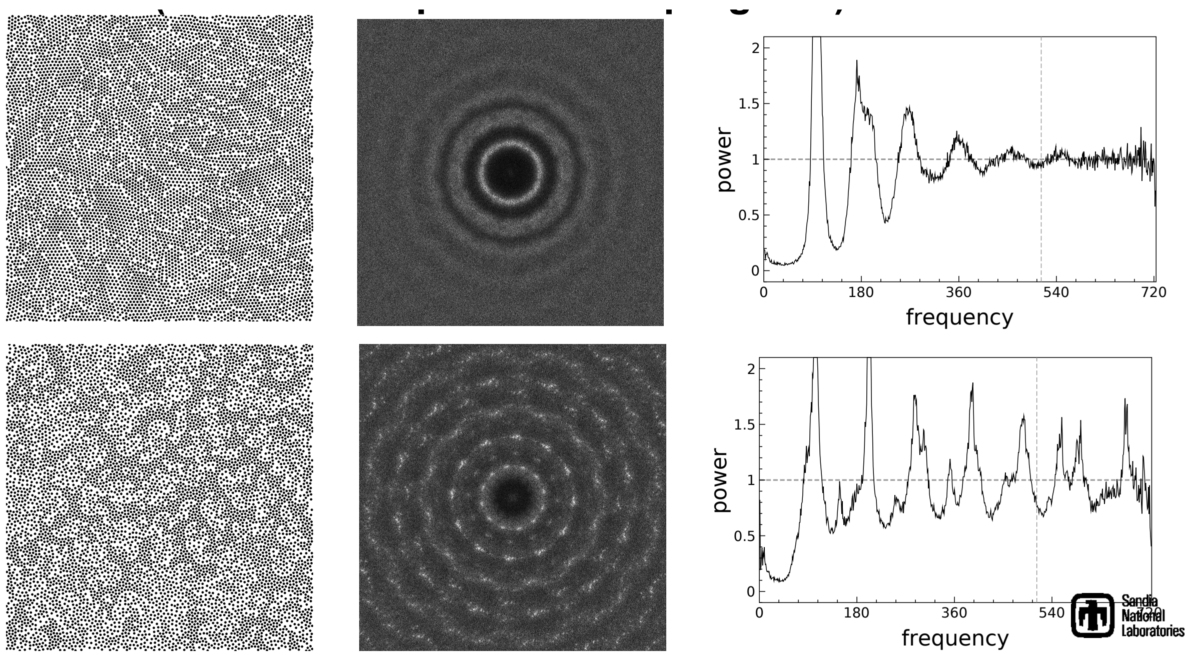

Figure 2: Delaunay refinement (Triangle) point clouds from particular choices of target edge lengths and angles, visualized using the PSA tool. Top left we see patches of hexagonal packings, and bottom left we see circular patterns of jumps in the point spacing. In the Fourier transform, middle column, in the top the rings are more pronounced than form MPS; in the bottom we see bright spots which indicate preferential directions, meaning nearby points are more dense in certain directions than others. In the radial average, right, on both the top and bottom we see accentuated spikes.

Bonus subproblem E: do these problems for dimensions other than 2. Three to five dimensions have some graphics applications.

Bonus subproblem F: characterize the effect of the domain boundary, for non-periodic domains.

Partners: Alexander Rand, Mohamed Ebeida, John Owens, Anjul Patney, Andrew Davidson, Chandrajit Bajaj.

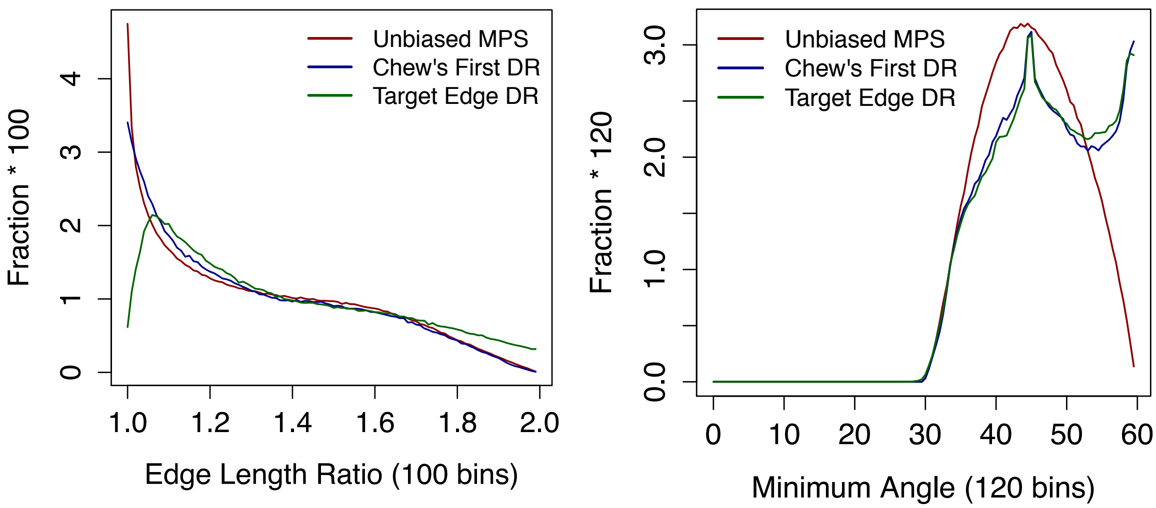

Figure 3: Computational Geometry measures of point clouds. Edge length and angle histograms of DR and MPS output. The minimum angle is the smallest angle of each triangle. The edge length ratio is the ratio of the length of Delaunay edges to the disc radius (MPS) or maximum Delaunay circumradius (DR). In both MPS and DR, the theories guarantee $r < |e| < 2r$.